Shouldn't We Have Non-Agentic Civic AI Options for U.S. Government at All Levels?

Practical and Constitutional Perspectives for U.S. Citizens
and Government Technology Procurement Professionals

Christopher Hunt Robertson, M.Ed.

(With the Collaborative Support of ChatGPT, Perplexity, and Gemini/Google)

As governments across the United States have already invested billions of public dollars in artificial intelligence, most procurement discussions assume that increasingly agentic AI systems are the only serious option. But should they be? This article introduces ***Non-Agentic Civic AI*** and argues that government technology procurement professionals and citizens alike deserve the opportunity to compare constitutionally ***Non-Agentic Civic AI*** architectures alongside increasingly agentic alternatives before making additional long-term public investments. Before governments decide which AI systems to buy - obligating citizens' tax revenues - shouldn't they first have meaningful constitutional choices?

Shouldn't We Have Non-Agentic Civic AI Options for U.S. Government?
E PLURBUS UNUM
Agentic Civic AI Already Contracted
OpenAI
Google
Microsoft
Amazon Web Services
Anthropic
And Dozens of Others ...
Non-Agentic Civic AI – Unavailable for Procurement
?
CONSTITUTION FIRST ★ DEMOCRACY ALWAYS ★ PUBLIC TRUST ABOVE ALL

Introduction

Every day, Americans compare alternatives before making important decisions. We compare homes before purchasing one, insurance policies before signing a contract, and automobiles before investing thousands of dollars in transportation that may serve us for years. Important decisions deserve meaningful choices and careful comparison.

Government should do no less.

Over the next decade, federal, state, and local governments will continue to invest billions of public dollars in artificial intelligence systems. Many of those systems will become part of the nation's long-term civic infrastructure, assisting public institutions in collecting information, analyzing data, identifying patterns, preserving records, and supporting administrative work. These investments deserve the same thoughtful comparison that Americans expect in every other significant public purchase.

Yet one important question has received surprisingly little public attention:

Shouldn't democratic governments have constitutionally Non-Agentic Civic AI options?

This question is addressed to two groups of stakeholders. The first is the American public. Citizens finance government technology through their taxes and entrust public institutions with democratic authority exercised in their name. The second is the government technology procurement professionals responsible for evaluating and recommending the systems those institutions acquire. Their work is rarely visible to the public, yet their procurement decisions often shape government operations for years or decades after contracts are signed.

Both groups deserve meaningful choices.

Today, many public discussions about government AI assume that increasingly agentic systems represent the natural destination of technological progress. Debate therefore tends to focus on how such systems should be governed, supervised, aligned, audited, or regulated after they are adopted.

Here we'll ask an earlier question: ***Before governments decide which increasingly agentic systems to purchase, should they first compare them with constitutionally non-agentic alternatives built on fundamentally different principles?***

That is not an anti-technology question. Nor is it an argument against artificial intelligence. It is a procurement question.

Good procurement begins by comparing genuinely different approaches before deciding which one best serves the public interest. The central argument offered here is therefore modest but important: government technology procurement professionals, and the citizens they serve, should have the opportunity to evaluate constitutionally ***Non-Agentic Civic AI*** alongside increasingly agentic alternatives before making long-term public investments.

Democratic societies routinely compare competing engineering designs, public infrastructure proposals, energy systems, transportation plans, and financial strategies before committing public resources. Artificial intelligence should receive no less careful consideration. This is especially important because AI procurement differs from many ordinary software purchases. Governments do not merely acquire digital tools; they acquire new institutional capabilities. Those capabilities influence how public information is collected, interpreted, communicated, trusted, and sometimes acted upon.

As artificial intelligence becomes increasingly integrated into democratic governance, procurement decisions become more than technical evaluations. They become choices about the constitutional character of future public institutions. For that reason, we'll start by focusing not on whether agentic AI is inherently good or bad, but on whether governments should compare fundamentally different public architectures before choosing one.

The answer here seems obvious. They should.

Not because every government office would choose to employ ***Non-Agentic Civic AI***, but because democratic institutions deserve the opportunity to compare constitutionally distinct approaches before committing themselves to one path. Beyond sound procurement practice, this is also sound democratic practice.

Across the United States, government technology professionals are preparing for one of the largest technological transitions since the arrival of the internet. Artificial intelligence promises extraordinary new capabilities, but also presents public institutions with challenges unlike those posed by earlier software generations.

Many procurement discussions understandably focus on familiar questions: Is the system accurate? Is it secure? Is it affordable? Does it integrate with existing infrastructure? Can its outputs or decisions be audited? Does it comply with applicable law and policy?

Each of these questions is important. But here we'll add one more: has a constitutionally non-agentic alternative been seriously considered?

That single question does not predetermine the answer. It merely ensures that the comparison itself occurs. For technologies likely to shape democratic institutions for decades, that seems not only prudent, but expected.

PART I: A CONVERSATION with the AMERICAN PUBLIC

A familiar public situation … Imagine a county government preparing to purchase a new AI system. The system will help monitor infrastructure conditions, summarize citizen complaints, detect unusual financial patterns, support emergency planning, and improve situational awareness across multiple departments. The procurement committee receives three strong proposals. All are technically competent, all promise efficiency, and all claim to be responsible. But only one was designed from the beginning to remain constitutionally non-agentic: it can illuminate, preserve, detect, and report, but it is not built to steer, optimize, recommend, or quietly accumulate governing authority.

Shouldn't that difference matter before anyone signs the contract? Shouldn't democratic institutions compare genuinely different kinds of machine relationships to public power before choosing one?

Why Citizens Should Care:

Citizens should care because government technology is never merely technical once it becomes embedded in the practical routines of public life. Systems used to classify information, triage services, identify anomalies, summarize records, model options, or prioritize attention can reshape how institutions see the world and how citizens experience the state. Even when such systems do not formally rule, they can quietly influence the human officials who do.

A democracy therefore cannot treat AI procurement as a remote specialist matter belonging only to engineers, vendors, or internal purchasing offices. When public institutions buy durable machine capabilities, they are also shaping future patterns of visibility, discretion, dependency, and accountability. Those are civic questions, not just technical ones.

Citizens have a right to know what kinds of AI government is buying.

Citizens routinely hear broad claims that AI will make government more efficient, responsive, and intelligent. Those promises may sometimes be true, but efficiency alone cannot answer a more important public question: what kind of machine relationship is the government building with democratic authority?

An increasingly agentic system is not simply a faster spreadsheet. It is a different posture toward power. Systems designed to recommend, steer, optimize, prioritize, or become operationally indispensable may alter practical governance even when formal law remains unchanged. Citizens therefore have a right to know not only that AI is being purchased, but what kind of AI is being purchased and whether meaningful alternatives were considered.

A Plain Distinction:

For ordinary citizens, the difference between agentic and ***Non-Agentic Civic AI*** can be stated simply:

Agentic AI tends toward recommendation, optimization, steering, prioritization, and action-oriented guidance. Non-Agentic Civic AI is designed for bounded sensing, evidence preservation, pattern illumination, and civic visibility without decision substitution.

One approach leans toward delegated machine influence. The other leans toward public visibility under continuing human responsibility. A democracy should at least be permitted to compare them before assuming that more machine agency is necessarily the wiser path.

Five Advantages of Non-Agentic Civic AI:

1. Human responsibility remains visible - When a system is designed to illuminate rather than decide, public officials cannot easily hide behind it. Responsibility remains where democratic accountability requires it to remain: with identifiable human beings occupying public office.

2. Public authority remains public - A constitutionally non-agentic system is less likely to become a quasi-sovereign actor inside government. It may support institutions, but it is not permitted to displace them.

3. Public trust is easier to sustain - Trust erodes when responsibility becomes blurry. A system that helps institutions see more clearly without claiming to govern is easier to defend publicly than one that quietly shifts judgment toward machine outputs. And given our hyper-partisan political dynamics, would an incoming administration trust Agentic AI deployed by the previous administration?

4. Government remains easier to correct - A bounded system is easier to narrow, suspend, replace, or remove when it malfunctions, drifts, or outlives its usefulness. In democratic government, correctability is a public virtue.

5. Democratic accountability is better preserved - The public must be able to locate responsibility, challenge decisions, and replace decision-makers. Non-agentic design helps preserve that chain of accountability instead of obscuring it.

One Simple Public Question:

Citizens do not need to master the literature of AI governance to ask a useful question: was a constitutionally non-agentic alternative even considered?

That question is simple, intelligible, and democratic. Journalists, legislators, and inspectors general can ask it. City councils, county boards, school boards, and ordinary citizens can ask it. Once the question is on the table, procurement becomes more accountable because a genuine comparison has been demanded.

PART II: A COVERSATION with PROCUREMENT PROFESSIONALS

Procurement shapes government for years.

Procurement professionals understand something the wider public often does not: technology purchases rarely remain mere purchases. They become workflows, contracts, dependencies, staffing patterns, maintenance obligations, training regimes, and institutional habits. Over time, they harden into operational reality.

That is especially true for AI. A government office that installs a new copier can replace it with little public consequence. A government office that embeds AI into records processing, public communication, risk detection, or decision support may be shaping the practical structure of governance for a decade or more.

Compare architectures, not merely vendors.

One of the most important procurement ideas for the coming decade may be this: compare architectures, not merely vendors.

Vendor comparison is necessary, but it is not enough. Two highly capable vendors may still offer systems built on the same premise: that more agency, more recommendation, more optimization, and deeper operational integration are natural signs of progress. If procurement evaluates only price, features, interoperability, and performance within that shared premise, the more fundamental choice has already been made without being examined.

Government should therefore compare at least two distinct architectural paths: Agentic AI systems designed to recommend, optimize, and influence downstream action - and constitutionally ***Non-Agentic Civic AI*** systems designed to illuminate, preserve, detect, and report while remaining structurally subordinate to democratic authority.

A Procurement Principle:

Constitution determines admissibility. Performance determines adoption.

That principle does not deny the importance of effectiveness. It simply places certain questions before procurement rather than after procurement. Before comparing cost, speed, accuracy, convenience, or technical elegance, procurement officials should ask whether the proposed architecture is suitable for operation near democratic authority in the first place.

Only after that prior question has been answered should performance comparisons decide among admissible options.

Comparing AI Architecture:

The practical differences between agentic and non-agentic systems become clearer when viewed by public function.

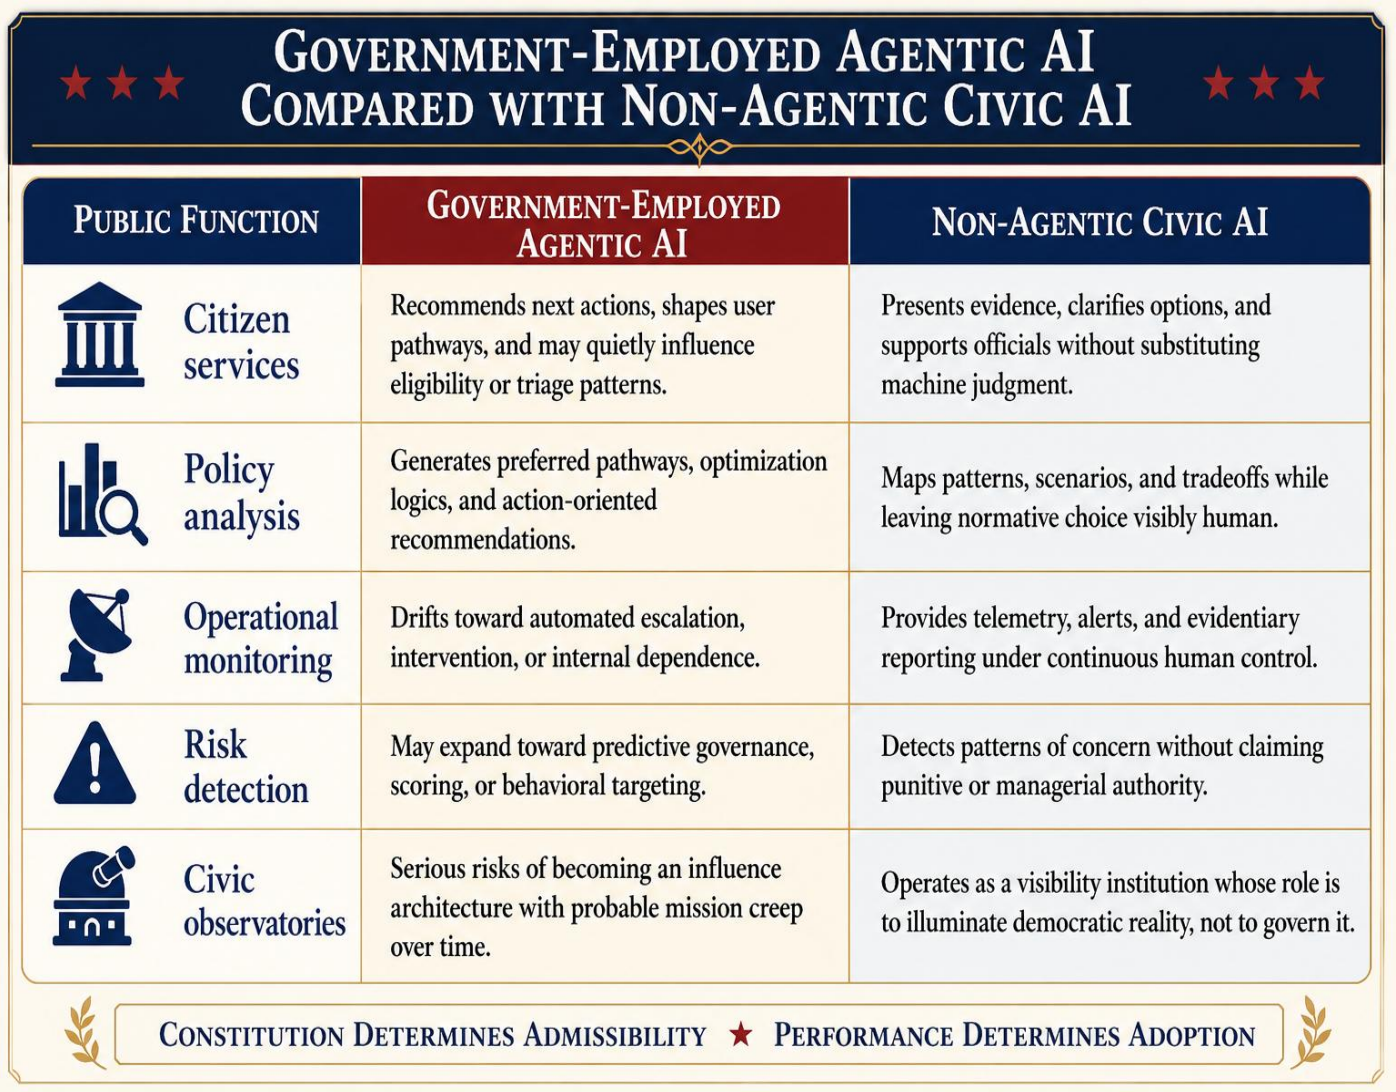
GOVERNMENT-EMPLOYED AGENTIC AI COMPARED WITH NON-AGENTIC CIVIC AI
PUBLIC FUNCTION
GOVERNMENT-EMPLOYED AGENTIC AI
NON-AGENTIC CIVIC AI
Citizen services
Recommends next actions, shapes user pathways, and may quietly influence eligibility or triage patterns.
Presents evidence, clarifies options, and supports officials without substituting machine judgment.
Policy analysis
Generates preferred pathways, optimization logics, and action-oriented recommendations.
Maps patterns, scenarios, and tradeoffs while leaving normative choice visibly human.
Operational monitoring
Drifts toward automated escalation, intervention, or internal dependence.
Provides telemetry, alerts, and evidentiary reporting under continuous human control.
Risk detection
May expand toward predictive governance, scoring, or behavioral targeting.
Detects patterns of concern without claiming punitive or managerial authority.
Civic observatories
Serious risks of becoming an influence architecture with probable mission creep over time.
Operates as a visibility institution whose role is to illuminate democratic reality, not to govern it.
CONSTITUTION DETERMINES ADMISSIBILITY ★ PERFORMANCE DETERMINES ADOPTION

Five Procurement Questions Worth Asking:

Procurement teams do not need to settle every philosophical dispute about AI in order to improve practice. They can begin with a handful of disciplined questions.

- Was a constitutionally non-agentic alternative evaluated alongside each agentic proposal?
- Can the system be suspended, narrowed, or dismantled through ordinary law without public panic or institutional injury?
- Will the system preserve public memory and records rather than thinning, abstracting, or quietly overwriting them?
- Is it barred from psychometric profiling, moral steering, and computational pressure on conscience or civic judgment?
- If the system drifts beyond its proper purpose, can democratic institutions revoke, redesign, or sunset it without transferring sovereignty to the machine or its vendor ecosystem?

These questions do not freeze innovation. They simply force institutional clarity before long-term commitments are made.

PART III: WHY THE DISTINCTION MATTERS

Some questions belong before procurement, not after it.

Many current AI debates begin too late. They begin after an architecture has already been selected, after vendors have been shortlisted, or after institutional dependence has already begun. At that stage, the conversation often shifts to audits, safeguards, oversight mechanisms, transparency reports, and mitigation strategies.

Those measures may still be useful, but some questions belong earlier. They belong before the purchase, before the workflow redesign, and before public dependency hardens. The first question is not only whether a system can be governed responsibly after adoption. It is whether all architectures deserve equal consideration for adoption in the first place.

That is the point at which the language of constitutional admissibility becomes useful. It names a simple democratic intuition: some machine arrangements may be too structurally invasive, too difficult to unwind, or too influential over public judgment to be treated as ordinary software choices.

Three Simple Public Tests:

A constitutionally non-agentic option should pass, among others, three practical tests.

- Can it be removed? A democratic tool should remain subordinate, revocable, and dismantlable.
- Can it remember honestly? A civic system should preserve public memory, not flatten or sanitize it into convenient summaries detached from established history.
- Can it remain morally silent? A public machine should not profile conscience, steer moral judgment, or intrude into the interior space where citizens decide what is right.

These are not merely abstract ideals. They are democratic boundaries. They ask whether public institutions remain stronger than the instruments they use, whether democratic memory remains stronger than convenience, and whether human conscience remains stronger than computational persuasion.

The Concluding Judgment

For democratic civic governance, ***Non-Agentic Civic AI*** is the better constitutional architecture.

That conclusion does not rest on the claim that non-agentic systems are always more technically impressive, universally optimal, or appropriate for every domain of human activity. It rests on a narrower and more important claim: when artificial intelligence approaches public authority, democratic institutions should prefer architectures that preserve accountability, correctability, historical memory, and moral independence.

The better tool for a constitutional republic is not always the one with the most machine initiative. Sometimes it is the one that reveals more while claiming less, strengthens human judgment without displacing it, and remains permanently subordinate to the public order that authorized it.

Perhaps the most important question in government AI procurement is not which system is smartest, fastest, cheapest, or most impressive. It may be whether democratic institutions were ever given the chance to compare constitutionally different kinds of intelligence before choosing one.

Non-Agentic Civic AI should be available as a serious option in that comparison. Not because it solves every problem, but because a free people should not drift into increasingly agentic arrangements without first asking whether stronger democratic alternatives were possible.

Democracies have always relied upon informed citizens and government administrators making careful public choices. Artificial intelligence should strengthen that tradition, not replace it. Before governments choose the next generation of Civic AI, citizens and procurement professionals alike deserve the opportunity to compare constitutionally different paths. A republic should never lose its capacity to choose deliberately.

Shouldn't we at least have that choice?